

Document made available under the Patent Cooperation Treaty (PCT)

International application number: PCT/GB05/000074

International filing date: 12 January 2005 (12.01.2005)

Document type: Certified copy of priority document

Document details: Country/Office: GB
Number: 0400569.0
Filing date: 12 January 2004 (12.01.2004)

Date of receipt at the International Bureau: 04 March 2005 (04.03.2005)

Remark: Priority document submitted or transmitted to the International Bureau in compliance with Rule 17.1(a) or (b)



World Intellectual Property Organization (WIPO) - Geneva, Switzerland
Organisation Mondiale de la Propriété Intellectuelle (OMPI) - Genève, Suisse



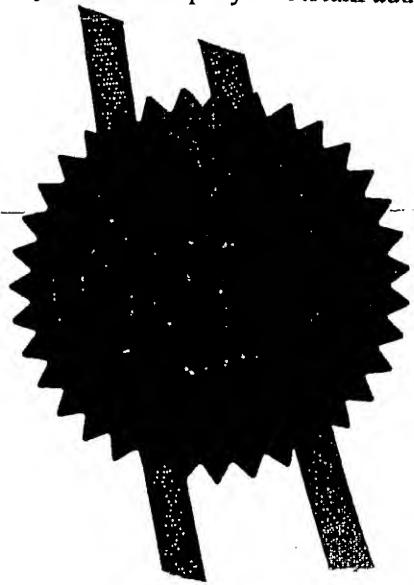
The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated 23 February 2005

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form.)



The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

AJR/47067.GB01

2. Patent application number

(The Patent Office will fill in this part)

0400569.0

3. Full name, address and postcode of the or of each applicant (underline all surnames)

Lightfoot Solutions Limited
Old Estate Office
Harleyford Estate
Henley Road
Marlow
Bucks
SL7 2DX

08786357001

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of incorporation

UNITED KINGDOM

4. Title of the invention

SYSTEM & METHODOLOGY FOR IDENTIFYING STATISTICALLY SIGNIFICANT EVENTS IN MONITORED PROCESS

5. Full name, address and postcode in the United Kingdom to which all correspondence relating to this form and translation should be sent

Reddie & Grose
16 Theobalds Road
LONDON
WC1X 8PL

Patents ADP number (if you know it)

91001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application
(if you know it)Date of filing
(day/month/year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day/month/year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.

See note (d))

YES

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form.
Do not count copies of the same document.

Continuation sheets of this form

Description	10
Claim(s)	3
Abstract	0
Drawing(s)	2 only <i>RM</i>

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (*Patents Form 7/77*)

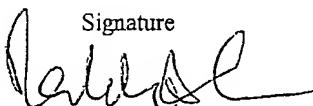
Request for preliminary examination and search (*Patents Form 9/77*)

Request for substantive examination (*Patents Form 10/77*)

Any other documents
(please specify)

11.

I/We request the grant of a patent on the basis of this application.


Signature

Date
12 January 2004

12. Name and daytime telephone number of person to contact in the United Kingdom

A J ROBSON
020-7242 0901

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or such direction has been revoked.

Notes

- a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- b) Write your answers in capital letters using black ink or you may type them.
- c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- d) If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- e) Once you have filled in the form you must remember to sign and date it.
- f) For details of the fee and ways to pay please contact the Patent Office.

Title: System & Methodology for identifying statistically significant events in monitored process

This invention relates to a method, apparatus, and computer program for identifying significant events in a monitored process.

Background to the invention

Ever since the concept of a process, or indeed any measurable activity, has existed there has been the desire to measure it and improve upon it. In the modern world we encounter processes everywhere from ordering a sandwich to the building of a car. Any of these processes can be graphically monitored by looking at some measure plotted against the relevant processes dimension, which is typically time, but could be any other dimension such as length.

In the 1950's a statistician called Deming took this idea to the manufacturing industry and showed that by applying a statistical rule to the data being displayed it was possible to show on the chart those points which were part of the normal variability of a process and those that were outside. These charts utilise three extra lines superimposed on top of the data, an average line and top and bottom process guidelines, the position of these lines on the chart being derived from the data rather than some arbitrary position. Significant events, by which we mean something out of the ordinary, are those points outside the process guidelines and by investigating and acting upon these; improvements in the performance of the process can be achieved. The charts are well known in the manufacturing industry as Statistical Process Control (SPC) charts and have been widely used in manufacturing since the early 1950's to great effect.

The problem faced by those seeking to implement SPC charts for a process within a business is that there is simply no enterprise wide, simple to configure, general purpose SPC tool that is relevant to every one in an organisation. To date, SPC charts and the software that displays them have remained in the domain of the statisticians and engineers looking after complex manufacturing processes.

Summary of the Invention

Embodiments of the present invention include a system method and computer program for obtaining measure data from one or more external systems, processing & analysing the data, storing the data, and then displaying the data in the form of a dashboard of dials which link to various charts, including SPC charts. Each person in the organisation has the option to view charts that are uniquely relevant to them. Furthermore, preferred embodiments introduce new statistical functions to the SPC charts so that business processes can be accurately modelled. Preferably the system uses standard web browser technology to view the charts and scales from a single user installation to the entire organisation.

One embodiment of the invention provides a method of storing organisational information separate from the measure data in such a way that an organisational hierarchy (OH) can be modified, including the addition of hierarchy levels, without change to the Database Schema or measure data, thereby producing a more maintainable database.

The OH is, in general, a set of hierarchies, joined into one hierarchy by grouping nodes, analogous to folders or directories in an Operating System. Each such hierarchy can be considered a 'Dimension' from an OLAP perspective. Dimensions can also occur anywhere in a hierarchy, to allow common structure at the top of hierarchies to be factored out. In the following example, dimension nodes are shown in bold, and the top-level grouping node in bold italic:

<i>Organisation</i>
Geographical
North
Engineers
Eng01
Eng02

```

...
Equipment
  Item1
  Item2
...
South
  Engineers
    Eng03
    Eng04
...
Equipment
  Item3
  Item4
...
Functional
  Sales
  Support

```

Here we see Engineers and Equipment dimensions nested inside a Geographical dimension to save repeating the North/South breakdown which is common to both.

Typically, Data is related (or 'belongs') only to the leaves of the hierarchies, though due to pre-aggregation or coarse granularity of data, this may not always be the case. Data is considered 'Appropriate' to a node in the hierarchy if it belongs to that node or any node below. Data Appropriate to a node is displayed in charts for that node. This method of aggregation allows the hierarchy to be restructured without any Data implications. For example, another level, Country, might be added in the above hierarchy below Geographical, with North and South now coming under each Country. Only the Database Table representing the Hierarchy itself would need to be modified, and then only its contents - no new columns or Dimension Tables would be required. A row of Data belongs to a node in this hierarchy, for a given Measure, if one of the ownership expressions in the definition of that Measure evaluates to the 'Owner ID' of that node. This is the only link between the OH and the Data.

This mechanism allows a row of Data to belong to different nodes in the OH for different Measures; the relationship between the Data and the OH is a parameter of each Measure Definition. This contrasts with the more common, and less flexible, Foreign Key link typical of OLAP systems.

Therefore, measure information is stored separate from the measure data so that the measure hierarchy can be modified without change to the measure data.

A method is also provided of jumping via a URL link from stored underlying measure data to a third party application.

Each Data Table used to hold the data from which Measure Values are calculated, can have one or more 'URL Templates' associated with it. Such a template is a mapping from values in a row of data in that table to a URL. Each URL Template is assigned to a particular column in the Data Table (typically one of the columns from which the URL is created).

If such a template exists for the Data Table used by a Measure, a corresponding option will become available on the menu from which a 'Data Form' is requested for that Measure. This Data Form will then show the values in the assigned column underlined and in blue, as is normally the case for Browser Hyperlinks. Clicking on a value in this column will open a new Browser window, and load it from the URL calculated by applying the URL Template to values from this row of data.

The mapping from Data Table values to URL is effected by means of 'placeholders'. A placeholder is simply the name of a Data Table column enclosed in braces. The following is an example URL Template:

<http://www.myCompany.com/myApp?ref={cRefNo}&emplId={cEmplId}>

(cRefNo and cEmplId are columns in the Data Table)

This Template would typically be assigned to the column cRefNo; if the cRefNo value has value 'R/1234' and the cEmplId in that row has value 'E567', clicking on the cRefNo Hyperlink will open a new Browser window loaded from URL:

<http://www.myCompany.com/myApp?ref=R%2f1234&emplId=E567>

(note that the '/' gets escaped as '%2f' as required in URLs.)

This allows integration with existing Web-based applications. An extension to URL Templates allows integration to Applications running on the client machine. If the URL Template begins with the string 'exec ' then the rest of the Template, once the placeholders have been replaced with the corresponding data values, is used as a local command string. In this case no escaping of special characters occurs. For example, the following would open notepad with the filename specified in the cFilename column of a Data Table, in folder C:\Temp (this would only work on a Windows client):

exec notepad "C:\Temp\{cFilename}"

This allows integration with locally installed and thick-client applications.

A method is provided for selecting data from multiple points in both an organisational hierarchy and a metric hierarchy, to allow disjunction in the case of multiple points selected within the same dimension, and multidimensional analysis in the case of points selected in different dimensions.

A method is provided for displaying a Correlation chart, which shows two or more measures displayed on the same chart together with guidelines that allow the degree of correlation to be assessed visually.

The correlation coefficient is an established statistic used to measure the degree of association between two data series. Thus given one series of data values:

$x_{t_1}, x_{t_2}, x_{t_3}, \dots, x_{t_n}$

observed at times

$t_1, t_2, t_3, \dots, t_n$

and a second series of data values

$y_{t_1}, y_{t_2}, y_{t_3}, \dots, y_{t_n}$

observed at the same times, the correlation coefficient may be calculated as

$$\rho_{x,y} = \frac{1}{n-1} \sum_{i=1}^n X_{t_i} Y_{t_i}$$

where

$$X_{t_i} = \frac{x_{t_i} - \hat{\mu}_x}{\hat{\sigma}_x} \quad i = 1, 2, 3, \dots, n$$

$$Y_{t_i} = \frac{y_{t_i} - \hat{\mu}_y}{\hat{\sigma}_y} \quad i = 1, 2, 3, \dots, n$$

$$\hat{\mu}_x = \frac{1}{n} \sum_{i=1}^n x_{t_i}$$

$$\hat{\mu}_y = \frac{1}{n} \sum_{i=1}^n y_{t_i}$$

$$\hat{\sigma}_x = \sqrt{\frac{1}{n-1} \sum (x_{t_i} - \hat{\mu}_x)^2}$$

and

$$\hat{\sigma}_y = \sqrt{\frac{1}{n-1} \sum (y_{t_i} - \hat{\mu}_y)^2}$$

The usual method that is employed to allow a visual assessment to be made of the degree of association between two data series is to construct a graph in which the data points have co-ordinates:

$$(x_{t_i}, y_{t_i}), \text{ for } i = 1, 2, 3, \dots, n.$$

However, instead of this, the Correlation Chart shows both data series plotted on one graph with common axes, one data series displayed as the co-ordinates

$$(t_i, X_{t_i}), \text{ for } i = 1, 2, 3, \dots, n,$$

the other data series displayed as the co-ordinates

$$(t_i, Y_{t_i}), \text{ for } i = 1, 2, 3, \dots, n,$$

so that time is the abscissa.

The benefits of this alternative approach are that:

(a) It preserves the order in time of the data, which is an essential property of performance data.

(b) The standard deviation $\hat{\sigma}_x$ is a measure of the dispersion of the data $x_{t_1}, x_{t_2}, x_{t_3}, \dots, x_{t_n}$, likewise the standard deviation $\hat{\sigma}_y$ is a measure of the dispersion of the data

$y_{t_1}, y_{t_2}, y_{t_3}, \dots, y_{t_n}$ Dividing by these standard deviations normalises the two data series so that they have similar dispersions and can be plotted on the same Correlation Chart with a common ordinate axis.

(c) The original two data series

$$x_{t_1}, x_{t_2}, x_{t_3}, \dots, x_{t_n}$$

and

$$y_{t_1}, y_{t_2}, y_{t_3}, \dots, y_{t_n}$$

will have a high degree of association, and hence be correlated, if a graph of the co-ordinates

$$(t_i, x_{t_i}), \text{ for } i = 1, 2, 3, \dots, n$$

and a graph of the co-ordinates

$$(t_i, y_{t_i}), i = 1, 2, 3, \dots, n,$$

display similar patterns, i.e. similar upward or downward trends or changes with time.

The transformations

$$X_{t_i} = \frac{x_{t_i} - \hat{\mu}_x}{\hat{\sigma}_x} \quad i = 1, 2, 3, \dots, n$$

and

$$Y_{t_i} = \frac{y_{t_i} - \hat{\mu}_y}{\hat{\sigma}_y} \quad i = 1, 2, 3, \dots, n$$

change only the location and scale of the graphs in the ordinate direction, so that graphs of

$$(t_i, X_{t_i}), i = 1, 2, 3, \dots, n$$

and

$$(t_i, Y_{t_i}), i = 1, 2, 3, \dots, n$$

will display identical patterns, and graphs of

$$(t_i, y_{t_i}), i = 1, 2, 3, \dots, n$$

and

$$(t_i, X_{t_i}), i = 1, 2, 3, \dots, n$$

will also display identical patterns. Hence if the original two data series display a high degree of association, so will the two transformed series, and *visa versa*.

The Correlation Chart contains two performance guidelines lines that are calculated from the data and used to assess the statistical significance of any correlation between the data series displayed on the chart. The usual approach that is used to assess the statistical significance of correlation is to test the null hypothesis that there is no correlation between two data series against the alternative hypothesis that the two data series are correlated. However, on a Correlation Chart, if two data series are highly correlated they will give co-ordinates that are close together and inside the region bounded by the performance guidelines. If data points fall outside the region bounded by the performance guidelines then that is taken to be a signal that the degree of association between the data series is weak, i.e. that they have low degree of correlation. In contrast to the usual approach then the performance guidelines in a Correlation Chart are thus a test of the null hypothesis that there is some specified, high, degree of correlation between the two data series against the alternative hypothesis that the degree of correlation is low.

The distance in the ordinate direction between the upper and lower performance guidelines in a Correlation Chart may be calculated as the range W of a sample of size two from a standard bivariate Normal probability distribution that will be exceeded with a small probability when the correlation coefficient ρ of the distribution is a specified, high, value. Numerical values of the range W are thus required in order that the Correlation Chart may be constructed. It may be demonstrated by a non-trivial mathematical argument that the required range may be calculated from the corresponding range W^* for a sample of size two from a standard bivariate Normal probability distribution when the correlation coefficient of the distribution is zero using the formula:

$$W = W^* \sqrt{1 - \rho}$$

Values of W^* are available in published statistical tables.

The above description of the Correlation Chart is presented in terms of just two data series. However, the arguments and formulae may be applied chart without alteration when the Correlation Chart is used to display three or more data series.

A method is provided for displaying a Regression Chart, which shows a measure plotted with the regression model calculated for that measure, so that goodness of fit can be visually determined.

The statistical method known as regression consists of fitting a model of the form:

$$y_{predicted, t_i} = \hat{\beta}_0 + \hat{\beta}_1 x_{1, t_i} + \hat{\beta}_2 x_{2, t_i} + \hat{\beta}_3 x_{3, t_i} + \dots + \hat{\beta}_p x_{p, t_i},$$

where

$$x_{1, t_i}, x_{2, t_i}, x_{3, t_i}, \dots, x_{p, t_i} \quad (i = 1, 2, 3, \dots, n)$$

are p series of data values ($p \geq 1$), known as "control variables" or "predictor variables" or "drivers", and recorded at times

$$t_i \quad (i = 1, 2, 3, \dots, n),$$

and

$$y_{predicted, t_i} \quad (i = 1, 2, 3, \dots, n)$$

are the values predicted by the model at the same series of times, and the parameters in the model

$$\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2, \dots, \hat{\beta}_p$$

are calculated from the observed values of the control variables and the observed values of the response variable

$$y_{observed, t_i} \quad (i = 1, 2, 3, \dots, n),$$

so as to minimise the sum of squares of the residuals

$$S = \sum_{i=1}^n (r_{t_i})^2,$$

where the residuals are calculated as

$$r_{t_i} = y_{observed, t_i} - y_{predicted, t_i} \quad (i = 1, 2, 3, \dots, n).$$

All the above constitutes established theory. The usual method that is employed to allow a visual assessment to be made of how well the model predicts the response variable is to construct a graph in which the data points have co-ordinates:

$$(x_{j, t_i}, r_{t_i}) \quad (i = 1, 2, 3, \dots, n)$$

for one or more of the control variables (i.e. for one or more values of j), or to construct a graph in which the data points have co-ordinates:

$$(y_{\text{observed}, t_i}, r_{t_i}) \quad (i = 1, 2, 3, \dots, n).$$

However, instead of this the Regression Chart shows the observed and predicted values of the response variable plotted on one graph with common axes, displayed as the co-ordinates:

$$(t_i, y_{\text{observed}, t_i}) \quad (i = 1, 2, 3, \dots, n)$$

and

$$(t_i, y_{\text{predicted}, t_i}) \quad (i = 1, 2, 3, \dots, n)$$

so that time is the abscissa.

Performance guidelines are calculated and displayed on the graph as lines $3 \times \sqrt{\frac{S}{n-p-1}}$

above and below the line representing the predicted values of the response variable.

When new data become available, for a time period after that used to calculate the parameters in the model, the model is used to calculate the predicted values of the response variable, and the observed and predicted values continue to be plotted on the graph together with the performance guidelines.

The benefits of this alternative approach are that:

- (a) It preserves the order in time of the data, which is an essential property of performance data.
- (b) The performance guidelines can be used to identify either individual times (i.e. positions on the time axis), or periods of time, where the observed values of the response variable differ significantly from the values predicted by the model.
- (c) The Regression Chart will show if the model remains valid for the time period after that used to calculate the parameters of the model – a breakdown in validity will be indicated by the occurrence of individual observed values of the response variable falling outside the region bounded by the performance guidelines, or by a run of observed values of the response variable all on the same side of the line representing the predicted values of the response variable.

A method is provided for modelling a process showing both a trend and cyclic variation with two or more seasons of data.

The method will be described for the case when the time partition is a month. The presence of cyclic variation then means that if a month gives a low result in one year then there is a tendency for that month to give low results every year, or if a month gives a high result in one year then there is a tendency for that month to give high results every year. The presence of a trend then means that there is a tendency for the data to increase steadily over the period of two or more years, or that there is a tendency for the data to decrease steadily over the period of two or more years.

When there are data for an exact number of years, the formulae needed to calculate the model and performance guidelines for the process are comparatively simple. However, if the results obtained are then applied to data for a period containing an incomplete year, the performance guidelines are unsatisfactory, particularly when there are only data for only two or three years.

To understand the reason for this, it is necessary to consider the correlations between individual data points and the model. Suppose that data for exactly two years, running from January one year to December the next year, are used to calculate the model. This means that the model for January is calculated as the average of the data for the two Januaries, adjusted for the trend. A data point for one of these Januaries will thus be highly correlated with the model for January.

On the other hand, the data point for January in the third year will be uncorrelated with the model, because it has not been used to calculate the model. In order that the same significance may be attached to any point that falls outside the region bounded by the performance guidelines, the distance between the line on an sChart representing the model and the lines representing the performance guidelines should take into account the degree of correlation between the data point and the model. If the degree of correlation varies from one data point to another, this distance should also vary.

However, it would be undesirable for the typical user of the sfn software to see the distance between the model and the performance guidelines varying, as they would have difficulty understanding the reason for this. The way to get around this difficulty is to adopt the principle that all the data are always used to calculate the model. This means that a data point will always be correlated with the corresponding model value, and it can be shown, by some non-trivial mathematics, that the same significance can then be attached to any point that falls outside the region bounded by the performance guidelines, to an acceptable level of approximation.

The formulae that are then needed to calculate the parameters of the model can be derived by some non-trivial mathematics and are as follows.

The method is based on fitting the model:

$$y_{k,j} = a_j + b \times i + e_{k,j} \quad \text{with} \quad i = q \times (k-1) + j$$

to the data points

$$y_{k,j}$$

where

$$j = 1, 2, 3, \dots, q$$

represents the months in a year ($q = 12$), and

$$k = 1, 2, 3, \dots, p$$

represents the years ($p \geq 2$).

The formulae apply when there are complete data for the first $p-1$ years, but data for only the first q' months for the last year, year p , and when $p \times q \geq 24$.

The parameters in the model are calculated as the values that cause the sum of squares to be minimised:

$$S = \sum_{j=1}^q \sum_{k=1}^p (y_{k,j} - a_j - b \times i)^2$$

In this summation, for the last year (when $k=p$) $j = 1, \dots, q'$ with $q' < q$.

For $j \leq q'$ we calculate a_j from:

$$y_{..J} = a_j + b \times \bar{i}_j$$

where

$$y_{..J} = \frac{1}{p} \sum_{k=1}^p y_{k,J}$$

$$\bar{i}_j = q \times (\bar{k} - 1) + j$$

$$\bar{k} = \frac{1}{p} \sum_{k=1}^p k$$

For $j > q'$ we calculate a_j from:

$$y'_{..j} = a_j + b \times \bar{i}'_j$$

where

$$y'_{..j} = \frac{1}{p} \sum_{k=1}^{p-1} y_{k,j}$$

$$\bar{i}'_j = q \times (\bar{k}' - 1) + j$$

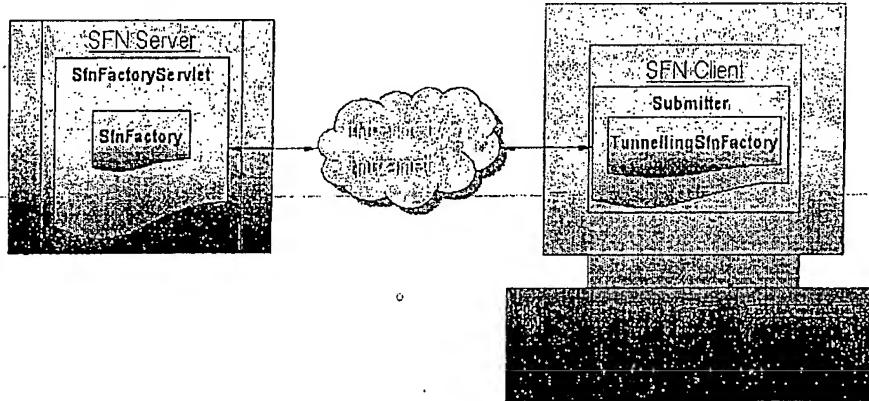
$$\bar{k}' = \frac{1}{p-1} \sum_{k=1}^{p-1} k$$

Thus if we know the trend parameter b we can calculate the parameters a_j representing the cyclic variation. The trend parameter b is calculated as:

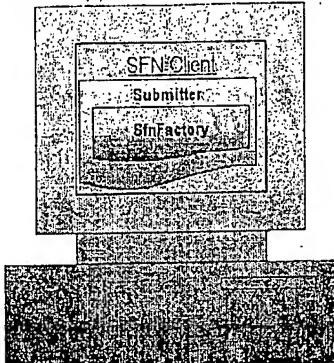
$$b = \frac{\frac{1}{q} \sum_{j=1}^{q'} \sum_{k=1}^p (y_{k,j} - y_{..j})(k - \bar{k}) + \sum_{j=q'+1}^q \sum_{k=1}^{p-1} (y_{k,j} - y'_{..j})(k - \bar{k}')}{\sum_{j=1}^{q'} \sum_{k=1}^p (k - \bar{k})^2 + \sum_{j=q'+1}^q \sum_{k=1}^{p-1} (k - \bar{k}')^2}$$

A method, that allows server client components of an embodiment of the invention to run together on a single machine without the need for an application server.

Interaction between the client and server components of the software occurs in the Submitter class on the client, via a single Java class which implements the interface SfnFactoryI. In the case of the Enterprise Version, this is the class TunnellingSfnFactory, which communicates with a Servlet, running on an Application Server (SfnFactoryServlet), over HTTP. The Servlet contains an instance of the class SfnFactory (which also implements the interface SfnFactoryI) which carries out all of the server-side functionality:



In the standalone version, everything runs on the client in a single Applet. In this case, the implementation of the SfnFactory1 interface used by the Submitter class is the SfnFactory class, obviating the need for an Application Server running a Servlet:



The Submitter class tries to load an instance of the SfnFactory class, to give standalone behaviour, but if this fails it loads an instance of TunnellingSfnFactory. This (required) failure in the Enterprise version is achieved by not supplying the SfnFactory class in the JAR file from which the Applet is loaded.

A method, that allows a user or session to override the default measure definitions by having a system of Layered Measure Profiles. A Measure Profile can be assigned to a User, and a 'Current' Measure Profile can be set for the session; these provide layers above the base Measure Definitions. Parameters that can be overridden include Display Options, Aggregation method, Data Filtering, and even the name of the Data Table used for the Measure (to allow pre-aggregated data to be used in some cases, for example).

Alternative Implementations

There are other software tools in the market place such as Minitab that display SPC charts from data input to them. These are tools for the statistician and differ from our invention in that these tools do not cater for an organisation (the organisation hierarchy), provide statistical functions applicable for modelling a business.

Disclosure date

Inventors

Mick Francis
David Anker
Roger Sym
Paul Turnbull

Claims

1. A method for taking process related data (**measure data**) from external system(s), processing that data, and storing it in a form that allows any individual to select and view an SPC chart that is of interest or uniquely relevant to them, comprising:
 - a. A database server,
 - implementing a schema that allows different measures to be defined and each and every measure value to be associated with a date/time stamp, and a user in an organisational hierarchy.
 - b. An application server running the server component of the invention that:
 - i. Presents to the user a mechanism for selecting where in an organisation hierarchy they wish measure values to be viewed from and what measure they would like to see.
 - ii. Provides a choice of different charts for the given selection including a dashboard of Dials that show the current value for each measure , a Benchmark chart that allows comparison of a selected measure across an organisational hierarchy, a Pareto chart that allows drill down through the measure levels and an sChart which is a form of SPC chart.
 - iii. Delivers the requested chart on demand to the user's Browser.
 - c. A web Browser connected via a computer network to the application server that loads the client component of the invention which interprets and displays the chart information sent from the server.
2. A method according to claim 1, that allows any user to select charts that are appropriate and permissible for them to view, comprising:
 - a. An **organisational hierarchy** presented in a tree format from which the user selects the point in the organisation from which charts are to be displayed subject to the user's permission level - See Fig1.
 - b. A **measure hierarchy** presented in a tree format from which the user selects the measure or measure attribute they wish to view - See Fig1.
 - c. A menu of functions applicable to the selected points in the hierarchies comprising:
 - i. A function to display a Dashboard of Dials for all measures for the selected point in the organisational hierarchy.
 - ii. A function to display an sChart
 - iii. A function to display a Benchmark chart
 - iv. A function to display a Pareto chart
 - v. A function to display the measure data in a tabular form.
3. A method according to claim 1, that shows the current value for each measure, comprising:
 - a. A **Dial** with highlighted segments that represent the normal variability of the measure data having values derived from the Performance Guidelines of the related sChart - see Fig2.
 - b. A pointer that represents the current value of the measure data.
 - c. A menu of functions associated with the measure and organisational point comprising:
 - i. A function to display an sChart
 - ii. A function to display a Benchmark chart
 - iii. A function to display a Pareto chart
 - iv. A function to display the measure data in a tabular form.
4. A method according to claim 1, that shows the comparison of a selected measure across an organisational hierarchy, comprising:
 - a. A **Benchmark chart** having on the X axis the names of each organisation which is a child of the selected point in the organisation , and on the Y axis, a scale suitable for displaying any measure value found in the measure in the organisations listed on the X axis.
 - b. For each organisation listed on the X axis, a vertical bar whose height and position on the chart represents the normal variability of the measure data having values derived from the Performance Guidelines on the related sChart - see Fig2.
 - c. On each vertical bar, marks indicating the average and the current value of the measure data for the particular organisation the vertical bar represents.

- d. A menu of functions associated with any particular vertical bar which the pointing device of the user's computer is currently over, comprising:
 - i. A function to display an sChart
 - ii. A function to display a Benchmark chart
 - iii. A function to display a Pareto chart
 - iv. A function to display the measure data in a tabular form.
- 5. A method according to claim 1, that allows drill down through the measure data, comprising:
 - a. A Pareto chart having on the X-axis the names of each attribute that is a child of the selected measure or attribute, and on the Y axis, a scale starting at zero and suitable for displaying any value found in the entries listed on the X axis.
 - b. For each entry on the X-axis, a vertical bar representing the sum of all measure values for that entry
 - c. A menu of functions associated with any particular vertical bar which the pointing device of the user's computer is currently over comprising:
 - i. A function to display a Pareto chart
 - ii. A function to display an sChart
 - iii. A function to display a Benchmark chart
 - vi. A function to display the measure data in a tabular form.
- 6. A method according to claim 1, that allows an SPC chart to be viewed for a particular point in the organisational hierarchy and for a particular measure or measure attribute, comprising:
 - a. An sChart having on the X-axis the dimension of the measure and on the Y-axis, a scale suitable for displaying any measure value found in the measure
 - b. For each measure value, a point drawn on the sChart whose X-axis value corresponds to the value's Date/Time, and whose Y-axis value corresponds to the measure value. Each point is connected via a line to the previous point and subsequent point bar the first and last points.
 - c. A horizontal average line calculated from the measure data.
 - d. Two horizontal Process Guidelines calculated from the average line, one set such that all points above the guideline are of statistical significance and one set such that any points below the guideline are of statistical significance. Statistical significance is calculated as any value outside of 2.66 times the average moving range from the average line.
 - e. A menu of functions associated with the sChart comprising:
 - i. A function to display a Pareto chart
 - ii. A function to display a Benchmark chart.
 - f. Context sensitive information attached to the point device tip (e.g. mouse cursor) comprising:
 - i. When over a point, display of measure value
 - ii. When under X-axis, display of measure value and details of any statistical process that has been applied to the measure data.
- 7. A method of alerting a user that there is a signal in the recent data of the sChart related to a Dial. The background of the dial changes to a different colour according to the following signals:
 - a. A run of points above the average
 - b. A run of points below the average
 - c. A single point above the upper Process Guideline
 - d. A single point below the lower Process Guideline.
- 8. A method of applying statistical functions to an SPC chart while it is displayed which allow behaviour in the real world to be modelled, comprising:
 - a. A function to insert a Process Break at a particular measure value. Calculation of the average and Process Guidelines restart from the selected measure value.
 - b. A function to model a cyclic process.
 - c. A function to model a trended process
 - d. A function to model a trended cyclic process
 - e. A function to exclude outlying data points
 - f. A function to limit the data range used in calculations
 - g. A function to annotate individual measure values.

9. A method of storing organisational information separate from the measure data so that the organisational hierarchy can be modified without change to the measure data.
10. A method of storing measure information separate from the measure data so that the measure hierarchy can be modified without change to the measure data.
11. A method of jumping via a URL link from the underlying measure data to a third party application.
12. A method of selecting data from multiple points in both the organisational hierarchy and the metric hierarchy, to allow disjunction in the case of multiple points selected within the same dimension, and multidimensional analysis in the case of points selected in different dimensions.
13. A method of displaying a Correlation chart, which shows two or more measures displayed on the same chart together with guidelines that allow the degree of correlation to be assessed visually.
14. A method of displaying a Regression chart, which shows a measure plotted with the regression model calculated for that measure, so that goodness of fit can be visually determined.
15. A method of modelling a cyclic process with only two seasons of data.
16. A method according to claim 1, that allows the server and client components of the invention to run together on a single machine without the need for an application server.
17. A method for allowing a user/session to override the default measure definitions.

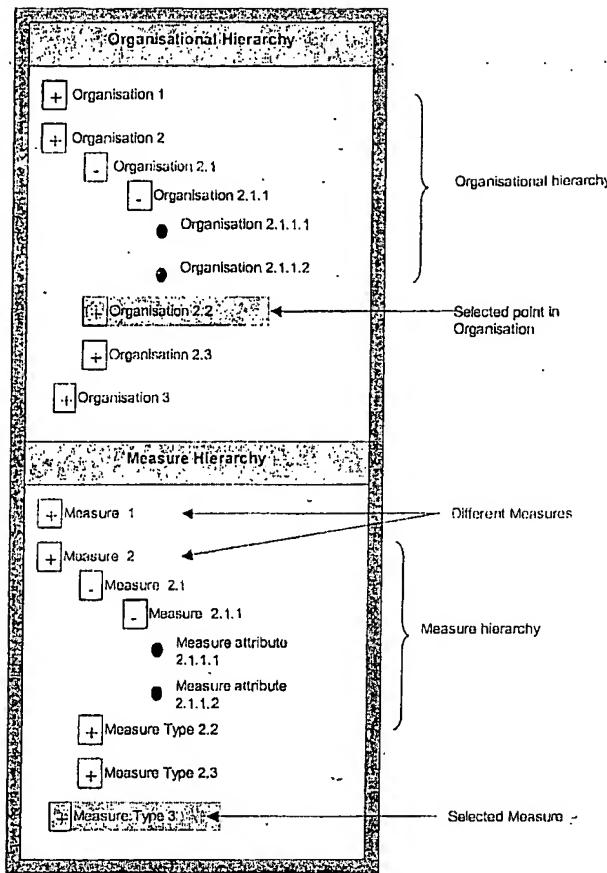


Fig1. Organisation and Measure Hierarchy's

Fig2 sChart

